NASA's Secured Airspace for Urban Air Mobility (UAM)
Virtual Workshop – November 1, 2022, 8 am to 2 pm (Pacific)
Hosted by the NASA Aeronautics Research Institute (NARI)

Table of Contents

Executive Summary

The planned UAM environment is leveraged from the Unmanned Traffic Management (UTM) concept of operations. There are various components like the Providers of Services for UAM (PSU), UAM operators, and Supplemental Data Service Providers (SDSP) which all provide services to support flight operations within that environment. The Federal Aviation Administration (FAA) can coordinate flight information between the FAA controlled National Airspace System (NAS) and the UAM environments through the FAA-Industry Data Exchange Protocol (FIDXP). This concept of UAM proposes to develop short-range, point-to-point transportation systems in metropolitan areas using vertical take-off and landing (VTOL) or short take-off and landing (STOL) aircraft to overcome increasing surface congestion. Understanding the various intricacies and communication exchanges, to garner the support of UAM and to realize its potential, an assurance of cybersecurity is critical for public acceptance.

To better understand how future UAM environments can be protected against cyber-attacks, and what mechanisms should be put in place to detect attacks against UAM environments - the NASA Secure Airspace Tech Group within NASA's Air Traffic Management Exploration (ATM-X), Urban Air Mobility (UAM) Subproject at Ames conducted a one-day Secured Airspace for UAM Workshop on November 1, 2022. This document represents an out brief of the discussions from workshop participants. The NASA Aeronautics Research Institute (NARI) at the NASA Ames Research Center organized the event. The goal of the workshop was to provide a setting for representatives within the UAM and cybersecurity community from various entities (government, private, and academic) to share insights into community needs, challenges, and solutions, and help to inform NASA's contributions to the UAM cybersecurity community.

Three major themes were focused on during the workshop either through presentations, anecdotal comments or through the question-and-answer portions throughout the workshop. These themes were identified as areas that were believed to be key for cybersecurity assurance of UAM. They included 1) identifying the challenges in Securing UAM Operations, 2) how to protect UAM operations and 3) how to detect those threats to the UAM environment.

The outcomes of the workshop are being shared with the broader community for consideration of challenges and opportunities to be used by NASA to consider means in which the agency might contribute expertise and technology to assist with cybersecurity efforts in the UAM space in the future.

Secured Airspace for Urban Air Mobility (UAM) Workshop Overview

NASA assembled a day-long workshop, this involved the introduction, identification of challenges in Securing UAM operation, understanding how to protect UAM operations, and

how to detect those threats to the UAM environment all broken down within the following nine (9) different sections:

To view the complete event webpage, visit: https://nari.arc.nasa.gov/securedairspaceuam
(Please reference Appendix A for the event agenda.)

Introduction

1. What is UAM - Annie Cheng (NASA ARC)

An understanding of UAM as a new market and overview of the FAA ConOps related to third party service providers. This will provide and introductory into ATM-X UAM airspace research and planned demonstrations for UAM airspace management (briefly mention that as we develop and test the architecture there is an increasing need for thinking about security).

Identify challenges in Securing UAM Operations

2. Challenges in Securing UAM Operations - Kenneth Freeman (NASA ARC)

Urban air mobility (UAM) is a concept that proposes to develop short-range, point-to-point transportation systems in metropolitan areas using vertical takeoff and landing (VTOL) or short takeoff and landing (STOL) aircraft to overcome increasing surface congestion. To realize the potential of UAM, an assurance of cybersecurity is critical for public acceptance. Cybersecurity has come to the forefront highlighting the need to protect these networks and systems from cyberattacks. The growth in the development of UAM systems, and the associated data exchange and service interactions will create abundant challenges due to numerous types of cybersecurity attacks. As these threats evolve, the UAM cybersecurity capabilities must adapt to these changes as well.

How to protect UAM operations

3. Principles of a Trustworthy System Gano B. Chatterji – Crown Consulting, Inc.

The expected Urban Air Mobility, Unmanned Aerial System, and Remotely Piloted Aircraft operations for passenger and cargo delivery will employ high degree of automation. Traffic management for these operations will be community-based and collaborative with operators coordinating, managing, and executing within the established legal framework. Data and information exchange between aircraft, operators, service/data providers, and regulators and public interest stakeholders such as law enforcement, cities, and safety organizations, will be accomplished with physical networks connecting geographically distributed sensor, communication, compute, and storage nodes. Data will be transported over wired data-

communication networks and wirelessly using air band radio and satellite communication between airborne, space-based, and ground-based systems. This cyber-physical world of interconnected and interdependent systems will provide numerous pathways for cyberattack to propagate to different parts of the larger system (system-of-systems) making it increasingly vulnerable. Cybersecurity challenges identified by government and industry include: (1) security and authentication for guarding against malicious activities and for preventing unlawful access to operator and government systems, (2) information/data integrity for ensuring information and data are conforming and trusted, and (3) data access for providing data to law enforcement, state and federal authorities enabled using remote ID information and correlation. This presentation introduces managers, engineers, system designers and programmers to ideas of risk, trust, and the key principles of engineering trustworthy systems. Approach to building trustworthy systems with a complete understanding of the system such as mission objectives, certification requirements, functionality, risks, design, architecture, and the policies and procedures for developing and operating the system is discussed. References to several cybersecurity frameworks are provided for additional information.

4. Security Architecture - Brian Nolan (NASA GRC)

This talk discussed using Model Based Systems Engineering (MBSE) for Systems Security Engineering. The ARMD Cybersecurity team has developed a process from a variety of sources; MBSE is being used as a facilitator of that process, to deal with the inherent complexity of analyzing large systems of systems.

Modeling threats, mitigations, and controls will be covered, as well as our work in progress trying to apply MIT's STPA-Sec to a functional architecture of the Advanced Air Mobility project.

5. Identity Access Management Panel - Adam Monsalve (CNA) & Robert Segers (FAA)

Identity and Access Management, or IAM, refers to the cybersecurity discipline concerned with ensuring the correct entities have access to the correct resources at the correct time. This involves verifying the identity that an entity is claiming to be in a digital transaction, the integrity of the transaction and verifying that the entity has the proper permissions to access the resource of interest. IAM considers not only the identity of a person entity accessing a resource, but also Non-Person Entities (e.g., servers, applications) exchanging data within the system. In the UAM ecosystem, there are numerous use cases that require IAM capabilities be in place, such as the securing of air-to-ground (e.g., Vehicle Command and Control) and air-to-air (e.g., Detect and Avoid) communications. A fundamental concept in enabling IAM is trust, and an ICAO effort known as the international aviation trust framework (IATF) will be an important contributor to a strong IAM capability for UAM. UAM stakeholders must consider the IAM needs as UAM concepts continue to mature to ensure that the ecosystem is resilient against threats that can be mitigated by IAM capabilities.

6. Blockchain for Aviation in UAM - Richard Walsh (NASA ARC)

The "Intermediate State" of NASAs Urban Air Mobility Maturity Scale, envisions operations with collaborative and responsible automated systems. The intermediate state is comprised of urban air mobility maturity level (UML) 3 and 4. UML-3, is characterized by low density, medium complexity operations with comprehensive safety assurance automation. UML-4 is characterized by medium density and complexity operations with collaborative and responsible automated systems. Securely transmitting Position, Navigation, and Timing (PNT) data is central to achieving this intermediate state. This concept explores one possible pathway toward achieving UML-3 and -4. Further, this concept explores the use of blockchain technology in an aviation operational environment to resolve securely transmitting time critical data e.g., Automatic Dependent Surveillance-Broadcast (ADS-B) in the open.

7. Immutable Secure Data Exchange and Storage for UAM - Kenneth Freeman (NASA ARC)

Urban Air Mobility (UAM) has become a focus for the next generation of aerial passenger transportation. UAM operations will be leveraging a service-based architecture for airspace solutions. The UAM environment will leverage diverse communications and system access approaches. These approaches include independent Providers of Services for UAM and supplemental data service providers that exchange data between themselves and UAM operators. This research focuses on the secure data exchange and storage of this decentralized UAM environment to address these challenges. This research intends to leverage a permissioned blockchain approach to address cybersecurity threats that may impact a UAM environment.

How to detect those threats to the UAM environment

8. In-Time National Airspace Security Event Identification using the In-Time Aviation Safety Management System Technology - Paul Nelson (NASA GRC)

Understanding the evolution of airspace operations and safety: can we predict evolving cybersecurity incidents in-time to mitigate safety impacts? Develop a clear understanding of the security posture for the system? And how do we continue to expand the computational modeling capability for security concepts? This talk aims to show how we can address these and meet the goal of developing a security analog to In-Time Aviation Safety Management System Technologies.

9. Model-based Approaches for Cyber Risk Assessment of Space Missions - Arun Vishwanathan (NASA JPL)

Space assets such as satellites, drones and spacecraft are increasingly attractive targets for adversaries ranging from individual hackers to hacker groups and nation states. This rapid rise

in the cyber threat landscape necessitates smarter cyber tools to stay ahead of our adversaries. In this talk, I will present the work being done by the Cyber Defense Engineering and Research (CDER) group at JPL in addressing cyber challenges as pertaining to JPL's missions. I will specifically focus on our work combining model-based engineering and AI-based reasoning to build a unique capability for automated cyber risk assessments.

Secured Airspace for Urban Air Mobility (UAM) Workshop

Speaker Introductions, Moderators: Terrence Lewis and Nishant Sharma, NASA Ames
Note Taker: Angela Boyle, KBR at NASA Ames

Welcome: Kevin Witzberger, NASA Ames, Sub-Project Manager for Urban Air Mobility Airspace (UAM)

- We have an effort in our Sub-Project using expertise in Secured Airspace: how to secure different data exchanges that we are exploring and being evolved with industry and the FAA during our research and development efforts.
- Super excited that about 275 people have registered for this particular workshop. What struck me the most is the diverse group of people from different organizations that I have seen across all of these NARI-hosted discussion topics. Thanks, NARI, for the service they provide. We have the FAA and NASA, of course, but so many people from other federal agencies, academia, a very wide cross-section of industry, state and local government representation. That tells me this is important.
- When we do research and development for airspace technologies, we tend to overlook the security aspect of it, and might address it later. This was an effort to bake the secured airspace part of it from Day 1, right in the Project. Ken Freeman, and others on his team, including Terrence Lewis are part of our Sub-Project, and help us to make sure we cover all our spots, and don't have our blinders on. It is usually two different domains that usually do not mix. The mixing took time, but very pleased with how it is working out. We are doing impactful work, but part of that is educating people like me, who didn't know about secured airspace in general, and become more educated and bake it into the research and development efforts that we have planned.
- I always learn something new when I am with Ken and team. Today, I will learn a lot of things I did not know. Hope you are all as excited in looking forward to this as me. It was intentional not to have this as a two-day event, or an eight-hour event. We did not want to overwhelm people. That is by design.
- Welcome to everybody. Thank you for registering and thank you for attending. This meeting is being recorded and will be available later.

Goals of the Workshop: Terrence Lewis, NASA Ames

- Work within the Secure Airspace Technology Group within the UAM Sub-Project.
- Over the course of this workshop, you will hear from speakers from three key areas:
    - Identify challenges in securing UAM operations
    - How we protect UAM operations
    - How we detect those threats to the UAM environment
- Shared the agenda and logistics. We will be using conferences.io feature. The link will be in the Teams Chat, and after each speaker there will be a Q&A portion. We encourage

community dialog in the Teams Chat but will not pull questions from Teams Chat due to time constraints.

- We have a great set of speakers, and we have some things we want to take away from the workshop. We will write up a report afterwards, and we will release that towards the end of the calendar year. The intent of this workshop is to make recommendations on how future UAM environments can be protected against cybersecurity attacks and what mechanisms should be put into place to detect those attacks. We ask that you keep those two things in mind during the course of the workshop to guide questions and thoughts as our presenters are speaking.

1. UAM

What is UAM? Annie Cheng, NASA Ames

- Nice to meet you all virtually. Good to see so many from different government agencies, as Kevin previously said.
- I work with Kevin as the Deputy Principal Engineer for the UAM Sub-Project.
- Do not need to go into too much detail on what is UAM but like to talk about the challenges we are facing, particularly in the airspace side of things.
- Shared an Advanced Air Mobility (AAM) diagram that covers a much broader concept than UAM. UAM is focused on a sub-set of that: passenger and cargo traffic within and around congested urban areas.
- Instead of fighting traffic on the ground, UAM looks upwards towards the sky to see if there are better ways to connect people from cities to different regions within the urban area and give them more possibilities to connect and travel.
- NASA is helping to move this industry along, how to leverage technologies in terms of the vehicle, and automated air traffic management. It is more than airspace, but today we will focus more on the airspace piece.
- It takes a lot to invest in this—what regulatory framework this will be working within, as well as how these new technologies are going to be developed.
- These are the use cases we are considering. The airspace folk usually don't talk with the security folk, but since working with Ken and team, we have learned so much more about connections that you need to think about up front, and you have to design with security in mind in order for these to work, especially in these environments.
- Part of what is interesting in UAM airspace is that we are looking at a notional architecture that came from the FAA's ConOps about two years ago. This was the starting point of what the end system is going to look like, and those who are familiar with **U**nmanned Aircraft Systems **T**raffic **M**anagement (UTM) this architecture came from UTM. There are good reasons why we started from that.
    - This architecture enables third parties, technologies developed by industry, that they could collaborate with each other, provide information as well as provide service to the operators. Some of the services are not going to be provided by the government, so we wait for industry to develop, implement, understand what the standards that need to be established for two vehicles to talk to each other.

    - A second characteristic is that we are relying more on digital communication and that is interesting because now we have a new challenge: how do you send information, such as aircraft telemetry information, operational intent of where these flights are going to go, how do you interact with the system actors? This array of technology that is going to come from industry, and the governance will require a centralized entity to provide governance to these different actors to make sure that they work collaboratively and fairly. This is being done with community-based rules (CBR) that industry is going to develop in order to talk to each other—rules of the road. These rules will talk to each other, and will be governed by some centralized authority, like the government. What will that picture look like, and how will we use it in the architecture?
    - As a Sub-Project we have what we call the X4 Simulation, just completed in June this year. We worked with 7 industry partners. This is part of the National Campaign simulation. We basically took the architecture and evolved it. We looked at what the provider of service to UAM is going to look like and developed requirements jointly with industry. We tested and implemented it as part of the effort. As part of the effort, we leveraged some of the UTM-developed technologies, such as FIMS, to authenticate and authorize the users. We also leveraged the DSS, how to synchronize information while doing so in a secure way. We have recently started some of those discussions. Today we hope to discover what other efforts we can look into, what other components of security we can try to implement and test in the future.
- This is a good segue to Ken Freeman, since we have been talking, there are a lot of connections that you don't see, but we are testing out what is the required data exchange, and what is the protocol that would be appropriate.

Questions from the conferences.io

- Q. Le Roy Ty sees public safety in there, and as an EMS helicopter pilot, and sees public safety in there. Regularly trauma centers are in the core of these urban areas. Will there be a new equipment/transponder that will interact with UTM that will be required by someone like me flying an EMS helicopter? We are regularly below 400' and regularly land on buildings
    - A. Annie Cheng: this is an interesting question, since you have the UAM traffic. We are working on where these airspaces would be, where UAM would operate within an airspace with associated airspace requirements. We don't have all the answers yet. Looking at existing aircraft and understanding how they will communicate is interesting to us.
- Q. The diagram does not include vertiports or other airport facilities. Is that deliberate? Or are these locations included elsewhere?
    - A. Annie Cheng: The diagram is software, system architecture, so it would not include any physical. But coordination of vertiports is something we are looking at too.
    - A. Kevin Witzberger – regarding the question on vertiports, the architecture diagram that Annie showed came right out of the FAA ConOps as a starting point and has

been evolving. This is not the final architecture. One of the evolutions not showing is the vertiport, so it is absolutely part of the architecture. It has software, components, including human components. It is just not shown in this version but is definitely included.

- Q. Since all drones (small, AM, etc.) will need to communicate across a single system, why separate conops for each?
    - A. Annie Cheng: A good question. We are learning that UAM, because of the low altitude it operates in, it does have interaction with existing traffic. That is one of the factors. There might be other factors. We don't necessarily want to leverage everything exactly as drones are doing.
- Q. The industry data exchange protocol, does that have a definition somewhere, or are there data types that are expected to be included in that?
    - A. Annie Cheng: Yes, the orange block. This came from the notional architecture from the FAA ConOps V1.0. So, there is an initial draft definition of what that includes. In terms of information exchange, it is still a research question. Information sent to get authorized, a concept of what "authorized" means in terms of different airspace classes, and different VFR and IFR rules.
- Q. To clarify, >5000 existing heliports are not towered, and private-use. How would these facilities interact with this architecture (as described by NASA's Automated Vertiport R&D project from last year?
    - A. Kevin Witzberger: this is a partial answer. Believe that the automated vertiport R&D Project last year, think this is a reference to NASA's Vertiport Project. That is a vision of automation at a future state, far term, high-density operations. There is an evolution to get there. So, part of the evolution is to figure out the role of the automation and the humans, both for towered and untowered airports/facilities. It is an active area of R&D.
- Q. Will the Airspace Authorization function on the left side of the diagram require human interaction or will it be an automated system?
    - A. Annie Cheng: This may best be answered by the FAA, but we are working with what authorization there is. In terms of information exchange, there might be some low-hanging fruit similar to how you would file a flight plan today, but if you are talking about clearance, it depends on where you are in the airspace, controlled, or not controlled. The beginning is an evolution. We are going to look at existing regulations and rules, so UAM will be operating in that environment, but slowly, and that is something we are looking into.
- Q. Are we looking at one federated system or multiple federated ones each of which with separate governance?
    - A. I do not have the answer. When we started with the architecture, the federated system, how do multiple PSUs talk to each other? Slowly we realized that a PSU may not be a single entity. Within that it is also federated. So, the problem becomes there might be regulations as to which entity of the federated system you are in, and that is still being expanded, so it does make you think that these boxes are not just boxes, but multiple boxes.

2. IDENTIFY

Challenges in Securing UAM Operations – Kenneth Freeman, NASA Ames

- I work with Kevin Witzberger, Annie Cheng, Terrence Lewis and others in the Secure Airspace area
- Will pivot this towards cybersecurity. Want to throw some definitions, then we will talk about threats – any circumstance or event that could harm a system or piece of infrastructure. What threats can exploit vulnerabilities. Threats may be intentional or unintentional they are really independent of the system itself. There could be a flaw in the software, caused by a design error, coding error, or implementation error.
- When we talk about risk it is a measure of something negatively impacting the system and is related to what is the independent threat to the system. What is the asset that we are focused on, and what is the vulnerability within that system that could be exploited by that threat that could provide an adverse impact to the system. We also could look at the perspective of the concurrence of that system.
- Annie talked earlier about the UAM environment. One of things we are looking at from a security perspective: is it a service-oriented architecture, organizations independent of the FAA, providing services within this environment?
- Those organizations will be independently providing either vehicle operations (UAM operations) or vertiport operations. Right now, in the architecture that is communicated through a Provider of Services for UAM (PSU) which is a system used for coordinating such things as airspace, access to airspace, those types of things. When we look at this from a security perspective, we are thinking of those PSUs and operators as independent functions, and data services supporting that infrastructure, like micro weather also would be independent entities as they move forward, and then there are the communications back to the FAA, and coordination with the FAA.
- I came from the NASA Security Operations Center and had cybersecurity as our starting point, and then identify and understand the organization, and understand the assets within that organization. Going back to the architecture diagram, what are the assets that support a data service? If I have a PSU, what are the assets within that?
    - The next piece would be the protection of that environment. There will be talks later on technologies looking at different elements of the environment.
    - The third piece is if something goes wrong, how do I know about it? How to detect and attack within an environment and then to assess whether that attack was successful.
    - The last are operational issues, and how to respond. In the UAM environment the question is across different organizations, and do you communicate that information or not.
    - We see the cybersecurity framework. The response to an incident is different from a recovery to an incident. We see instances where a system has been compromised and may be taken offline. Those are recovery aspects.
    - This comes from the National Institute of Standards and Technology – on the federal side, this is a common model used for cybersecurity.

- Some of the challenges:
    - This is a decentralized UAM operations which is supporting flight – we need to determine the overall cybersecurity governance policy across a decentralized environment that is operated by multiple entities. Questions are: will there be a common guidance on how we do vulnerability management? Within the federal government, that is in place. There is common guidance.
    - Once attacks happen, will there be guidance on whether we should share that information across multiple entities, e.g., if PSU is attacked with an attempt to exploit an API, should that be shared with another PSU operator?
    - Will there be a coordinated incident response plan?
    - There will be a lot of challenges on how these elements fit together.
    - From an operations perspective, there needs to be a trust model in terms of the communications between the different operators and service providers within a model. Trust models will be in a later talk today.
    - We also need to look at identity management, also will be discussed later.
- When we looked at the UAM environment, we started to break it up into pieces, not a singular system, there is a cyber-physical element—the vehicles and the vertiports. There are the people who are operating the systems, and there is the infrastructure. We need to think about cybersecurity from those perspectives as well as the functions of the systems and the data exchanges.
- Some of the threats that we see would be
    - Jamming leading to limitation of data, which would limit vehicle communications
    - Denial of service—we need some level of redundancy to protect against this
- Looking back at the challenges: GPS spoofing is one of the jamming threats. Encryption helps, and also alternative communications will help.
- Additional threats are social engineering. Phishing is the primary one, attempting to seek users in order to give away sensitive information. These are attacks on people. Ransomware has been around for about 10 years and is almost its own element of cybersecurity. History shows that if you pay, you will likely have another attack within six weeks.
- There is a broad range of threats. There is voice, where a person is called and asked to give away their certificates. These are also coming in by text, and then there is email phishing. There is also insecure design and system configuration. There is insider threat, either by mistake or intentional.
    - Some of the mitigations are education, zero trust, continuous monitoring, response and recovery plans
- Attacks on the infrastructure of the software elements are rising. Access control problems need to be looked at. The concept of a network of computers that have been highjacked can attack something else. The pace can multiply
- Other aspects: access control, cryptographic failures, vulnerable and outdated components within a system, maybe identity. One of the mitigations is to have a pre-incident strategy including backing up systems, asset management, restricted user privileges. Also, a post-incident plan: how to rebuild a system, whether to take a system offline, how to communicate. Strength of identity to expand multi-factor authentication.

- Decentralized operations is a challenge for security, Also the protection of cyber-physical systems, the operational technology elements in the IT community, human error, vulnerabilities within the software

Questions from the conferences.io

- Q. If an operator's system were compromised, would that compromise the AM system as a whole? Is there a "circuit breaker"?
  - A. Ken Freeman: there should be a circuit breaker. Things like networks where you could compromise a piece of the system, not the entire system. Identity access management is another circuit breaker. Once a compromise is detected, then a risk assessment has to be done.
- Q. I am a systems integration contractor with FAA UAS SysOps Security Integration Team. MY QUESTION: Does the NASA software vision address or integrate in blockchain to detect and mitigate falsely transmitted information to AM aircraft? I ask this because I believe the end result that the UM and (overarching AM) communities desire is that their UAS and AM vehicles would be able to detect a transmitted information "lie" by flagging identifying and ignoring inconsistencies within blockchain information trail.
  - A. Ken Freeman: will be giving a blockchain discussion later, as well as Richard Walsh, so from a vision perspective we are looking at how we can apply blockchain technology within the UAM environment, particularly mitigating falsely transmitted information. That is a focus on the research ideas. Agree that it could help with detection. Both of our talks later will focus on mitigation. Detection should be added
- Q. Do AM corridors mitigate cybersecurity issues (air-gapped networks, decentralized management, no cascading effects, etc.)?
  - A. Ken Freeman: Do not know if the corridors themselves help us from a cybersecurity perspective. They do help us from an operations perspective for sure. The decentralized management, would consider this a challenge, and is as weak as the weakest link. This could be problematic in a decentralized environment if we don’t have a common model of what standards to keep the security levels at.
- Q. L1/L2 GPS signal spoofing and jamming is relatively cheap and easy. When I was young, we used to load encryption codes into the GPS of our Apache helicopters. What mitigations are being discussed as it applies to AM architecture?
  - A. Ken Freeman: We are discussing redundancy internally with one of our sub-teams. The other blockchain talk coming up has a novel idea using blockchain tokens as part of that process. Redundancy is one of the key issues, and also how do we detect these types of attacks?
- Q. What is NASA considering for detecting a Cyber event on the partners infrastructure or over the air and how would that be communicated throughout?
  - A. Ken Freeman: We do not know. From a partner or a service provider that is coming within the environment, would assume that since NASA is a research organization, the FAA would be the operational organization for this. Do not know if the FAA or any other organization will have the authority to monitor for cybersecurity attacks within the infrastructure of the service partner. This goes back to the governance question. Will there be standards/ governance/ policies/

regulations that you have to provide information. Cannot speak for the FAA but cannot imagine that one side of the government will be monitoring. We need to think about this.

3. PROTECT

Principles of a Trustworthy System, Gano Chatterji, Crown Consulting

- Will talk about how to engineer a trustworthy system
- Outline
    - Motivation
    - Cyber Physical System Risk Model
    - Establishing Trust
    - Approach for Building Cyber Resilient Systems
    - Key Principles of Cyber Security Engineering
    - System Theoretic Process Analysis
    - Aerospace Systems Standards
- Motivation
    - The future ATM needs to ensure availability, integrity, confidentiality, and safety of operations
    - Safety of vehicles and operations is paramount
    - Security is critical, due to sensors, networks, and computer being more vulnerable to bad actors
    - Goal is to design and develop cyber-resilient systems
    - All aspects of the entire system need to be addressed
- Cyber Physical System Risk Model showing Threat Agent > Threat Action > Vulnerability > Adverse Impact produces Risk.
    - Common Threats: data interception, jamming, denial of service, masquerade, replay, software threats, supply chain
- Establishing Trust:
    - What is Trust: The foundation of Cyber-Resilient Autonomy. It is confidence that the system has confidentiality (blocks access without proper credentials), integrity (protects data from getting corrupted), availability (continues to operate even when attacked), and safety (continues to operate safely and protect crew and equipment in degraded conditions.
    - Used as an example, Apollo 13, which brought the crew safely home.
    - Needed is a comprehensive understanding of the system, and implement policies and procedures, e.g., hierarchy of mission objectives defined and documented, processes for tracing system requirements and functions, GUI and other software tools, coding standards, and tools to manage software development life cycle processes. Examples also include automated workflow processes, policies for chain of control, data collection, architecture enhancements, understanding failure modes, contingency for graceful degradation, data at rest policies.
- Approach for Building Cyber Resilient Systems

    - Shared diagram showing overlapping circles representing Architecture, Functionality and Cyber Resilient Systems
- Key Principles of Cybersecurity Engineering
    - 1.Cybersecurity's goal is to optimize mission effectiveness; cybersecurity is never an end unto itself.
    - 2. Cybersecurity is about understanding, and mitigating cyberattack risk.
    - 3. Assume your adversary knows your mission and cybersecurity system better than you; the opposite assumption is folly.
    - 4. Defense in depth without defense in breadth is useless; breadth without depth is weak.
    - 5. Failing to plan for cybersecurity failure, guarantees catastrophic failure.
    - 6. Cybersecurity strategy and tactics knowledge comes from deeply analyzing cyberattack encounters.
- Principle 1: Cybersecurity's goal is to optimize mission effectiveness; cybersecurity is never an end unto itself
    - It is about understanding and mitigating cyber-attack risk. Assume your adversary knows your mission and your cybersecurity system better than you. Failing to plan guarantees a catastrophic failure. Strategy comes from deeply analyzing cyber-attack encounters and a deep understanding of your system.
    - You have to worry about all aspects of the system: the architecture, the functionality, and the development process. There are overlapping considerations such as operations and the human role.
- Principle 2: Cybersecurity is about understanding and mitigating cyber-attack risk.
    - To build the system, there are the mission requirements, and concept of operations. You need to understand the certification requirements, human and automation functions. You understand the threats, the risks and the priorities. You architect the system for resiliency and mitigation. You worry about the software standards and the life cycle processes, build, test, deploy. You need to continuously test, even after building, and make sure every component of the system is functionally as designed, and according to mission objectives
- Principle 3: Assume your adversary knows your mission and cybersecurity system better than you
    - Do not build secrets into the system
- Principle 4: Defense in depth without defense in breadth is useless; breadth without depth is weak
    - Discussed layering cybersecurity approaches (people, technology, process)
- Principle 5: Failing to plan for cybersecurity failure guarantees catastrophic failure
    - System failures are inevitable
- Principle 6: Cybersecurity strategy and tactics knowledge comes from deeply analyzing cyberattack encounters
    - Good cybersecurity operations are as important as good design
- The System Theoretic Process Analysis (STPA) consists of four steps:
    - Define purpose of analysis

    - Model the control structure
    - Identify unsafe control actions
    - Identify loss scenarios
- Scoring by Risk Cubes
    - Described the ranking: probability and consequence
    - How to objectively measure risks is necessary
- Aerospace System Standards
    - These are the systems standards – please look at these if you are interested.
- Summary: Building a cybersecurity or cyber resilient system is not cheap.
    - If you are heading towards a catastrophic failure, should you worry about cybersecurity? Or should you worry about trying to protect the mission?

Questions from the conferences.io

- Q. Is there any exploration on NASA's side of the Blue Architecture ideas being promulgated by DOD as part of the Blue UAS program?
    - A. Gano Chatterji: This workshop is to educate the community and look at all of these frameworks so we can evolve the system
- Q. Will charts be made available to workshop attendees
    - A. Gano Chatterji: Yes. The meeting is being recorded, and the video and presentation charts will be posted to the NARI workshop website within a few days

Security Architecture, Brian Nolan, NASA Glenn

- Primary area of expertise is modeling. Have recently come to the cybersecurity area. Will talk about how we are using modeling to facilitate security engineering
- Paul Nelson is in the background to assist with any questions
- Concluding slide is the starting point
    - The desire to bake things in first resonated with me, so security and resilience will be emergent properties of the system. This is much cheaper than trying to patch it on later.
    - We are using Model-Based System Engineering (MBSE) to facilitate our SE. Currently working with the Advanced Air Mobility (AAM) Project
- We do modeling for two or three major reasons: communicate more effectively, manage complexity, work with data behind the models
    - Shared a “Fly Safely” model sequence diagram. Shows high level to begin with and can show or hide detail. This represents flying safely, sharing data, and communicating.
    - We looked at “what could go wrong”. Things can go wrong on their own, or by maleficent intent.
    - We look at rudimentary things first. This document was taken from NASA’s In-time System-wide Safety Management System (IASMS) ConOps.
    - If you are in flight, there are a lot of things that can trigger an off-nominal state. We need to consider if they are created by natural causes, or ill-intent. Safety and security are related.

- The model helps you to think, and shared examples of how you could disrupt the normal safe operation of the aircraft.
    - We have a Cybersecurity Management Plan and are building a cybersecurity model.
    - We are using the NIST 800-160 Volumes 1 and 2, Systems Security Engineering, which is an extension of the ISO 15288 manual. We are also looking at STPA-SEC
    - If we have a system of interest and we want to model the architecture, looking at its behaviors and structures, we want to apply to the model an attack model, which would be threat data, threat intelligence. As we put these together and do threat modeling, then we apply it to the system of interest to get information on how threats are realized and how they impact the system of interest. This feeds back into the system of interest architectural model and can also feed into our counter-mitigation design activity.
    - A different view shows how we can apply a threat model with iterative loops planned into the process. When we do threat modeling, we know there will be mission impacts. These will prompt us to create protection needs with security and safety requirements. We will continue to do this until we reach a level of refinement to apply to the system design
- We have a candidate threat model for the NAS and will likely apply some of this to our AAM model. We took part from the UTM model done several years ago, and we have started to put it into a modeling tool, into MagicDraw models.
    - A threat is a threat source that is an actor that executes an attack on a target or an asset that results in a consequence or impact to a mission.
    - We take the threat models, and apply them to an architectural model, and we start to look at protection needs.
- An overview was shared of the model that Brian Nolan has been developing. It is in MagicDraw. There is a tree, and there are various packages, e.g., for threats. The relationships explicit in the model are there, and you can drill down to the various pieces. All of the material is retained and is usable.
    - A taxonomy is created.
    - A set of threats can be applied to a set of assets, and we can derive a set of protection needs. We are dealing with pieces of data that are manipulatable, accessible, and displayable.
    - We are trying to manage the complexity
- The AAM architecture is tenuous and developing over time. We are trying to reason about classes of threats and protection needs.
- Other tools and processes are included, and mathematical models will likely be used in the future.
    - We are trying to improve our security systems engineering to help manage the inherent complexity by divide and conquer.
    - We are designing is so it is not an emergent property of the system.

Question from conferences. Io

- Q. What value will modeling give us?

    - A. Brian Nolan: modeling can help you manage complexity, and communicate more effectively. Content, for example, on spreadsheets is not a unified set of data. We are bringing material into a model, so it can serve as a central source of truth. The data is key, not necessarily the images
- Q. Are you able to model threats against interfaces between system components or actors, or just with the actors/components themselves?
    - A. Brian Nolan: I can model threats against any entity in the model—actors, interfaces, connections.
- Q. How will normalization/baselining occur with an influx of feeders and signals which may overwhelm the system for real time analysis, part controlled vs. uncontrolled airspace by partners which may differ and be consistent?
    - A. Brian Nolan: do not know offhand.

Break – 10:05 to 10:35 am Pacific

Identity and Access Management (IAM), Adam Monsalve, CNA, and Robert (Rob) Segers, FAA

- Rob Segers introduced himself. He is the Information System Security Architect for the National Airspace System at the FAA NextGen Office
- Adam Monsalve from CNA – We are focusing on the topic of identity and Access Management, and will give some fundamentals: what is IAM and why should it matter in cybersecurity? Relationship to Urban Air Mobility. Then will go into topics that relate to IAM: on the international front, and the topic of trust
- We have an overarching topic related to IAM, whether federated or centralized architecture is the most appropriate for UAM
- Believe we were going to have a poll on this topic.
- What is identity and access management?
    - It is a security discipline which ensures that only the correct entities have access to the correct resources at the correct times. Most of us logging into a computer are interfacing with IAM
    - Key concepts are Identification, authentication, authorization
    - In UAM, there will be multiple resources and users, who will require access to certain elements of UAM so that all resources are adequately protected
    - IAM applies to both persons and non-persons
- Added value of zero trust is that you are using the observable state
- What is IAM protecting? We have three buckets of organizations: government, commercial, public safety. Each might have different concerns when interfacing with UAM across the ecosystem.
    - Rob pointed out that the fundamental protection of information is really important.
- The International Aviation Trust Framework (IATF) was explained by Rob. The concept was developed by ICAO, and the Trust Framework was a study group, which is turning into an ICAO panel. It is a combination of an identity federation concept to create interoperable identities and protections.

- When you look at the airspace system at a global level, it spans every facet of aviation, from airline, aircraft, air navigation service provider, fuel supplier, caterer. Every operational interaction, from gate to gate is touched by this framework.
- Today, ICAO is a trust framework so the 193 countries that are a member of it develop standards and apply those standards to member states, which codify those into regulations used by aircraft manufacturers, operators to secure and provide safety for aviation.
- In a connected world where we have cyber threats, you have to go further than the paper-based trust framework, to an operational framework that assures the interoperability of information protection to compatibility of identities. The appropriate legal agreements need to be in place between the participants, as well as a set of agreed upon technical requirements. We have developed in ICAO a set of performance requirements for confidentiality, integrity and availability based on ISO27001 and NIST 800-53 v.4. This is for the information exchange. For the identity, we developed a common policy that can be applied so participants in the framework can have the confidence that the information is protected at a certain level and that it is compatible with the safety risk when they use that information. We believe that the trust framework, from an operational standpoint, needs to provide an identity federation. There is a trust anchor at the ICAO level that allows trust to be determined between entities that participate in the trust framework that have their own identity system.
- Trust but Verify is the auditing component. Looking here at adherence to the technical requirements, that information is coming from the source, and is protected.
- Adam has a question on identity federation: that it is Core 1 of the identity federation. Do you have an example of this? Rob: a simple example would be datacom – an aircraft flying from one nation to another. The air navigation service provider wishes to connect to the aircraft. The aircraft will be from one country which has its own identity system, and the ANSP is from another nation. A TLS or DTLS tunnel can be set up between two entities, but mutual identification must be shared. How is trust established? This is where the federation comes in to provide the authentication. There are architectural principles on how to create the trust. A clean tree structure must be adhered to that always goes up to a pinnacle. If you allow horizontal trust relationships, you create the "ball of yarn". Adam: The Federal Bridge is a US overarching system that link your identity to a PIV card. It became a victim of its own success.

- Adam will take this topic. We want to give some examples where threats may be mitigated by IAM in UAM. Spoofing, Data Leakage, and Denial of Service
  - Spoofing by a participant giving itself an advantage is entirely possible and could be detrimental to competitors. IAM will be critical to prevent spoofing
    - Rob – the ability to spoof a UAM or a UAS, and you may have a malicious intent. Law enforcement can only detect good and bad, but spoofing can delay that determination and law reaction. Also, spoofing airspace is possible

    - Data Leakage – Adam there is a lot of sensitive data in UAM. The more we understand the roles in UAM, the better we can tightly restrict access and improve the security posture of the ecosystem
        - Rob: a successful attack is always preceded by a gathering of information. Then they will use that information to stage an attack
    - Denial of Service – Rob: this is a mixed bag and can lead to additional work to be able to authenticate information. The attacker could send false messages which you would have to triage, and may overwhelm you. Best method is to use standard practice first, and then have an efficient way to do your authentication, so the useful messages can be authenticated.
        - Adam: Denial of service in UAM is extremely important..Rob: you can pivot so that you are basically creating a side channel and the attacker gets completely cut off.
- Major Topic: IAM in UAM: Federated or Centralized?
    - Adam: How are we managing identities broadly across the ecosystem. Maybe ANSP, that ensures that IAM is managed in the same way uniformly across the ecosystem. Then there is the federated IAM. Rob touched on earlier, where you segment the identities into groupings or regions, and each region is responsible for its own identity management. There are pros and cons to both of these.
    - Rob: my philosophy is that it is not one size fits all. You really need to have a federation of centralized identity management, and the level of centralization is really left to the user community, but even with this, at some point in time, you will have to go outside this community, and you will end up federating your centralized identity system with others. By using a clean tree structure, you are able to manage your scope of trust. If you think of the multi-level tree, depending on which branch you put your trust anchor, you can have a de facto limitation of which identities you accept in that community. As you move up in the tree, your trust becomes broader. You can change and manage that scope of trust. Also, do not want people to confuse societal trust with technical trust or confidence. With trust, we are usually talking about confidence and technical trust, based on a set of agreed requirements. Societal trusts are up to the relying party and the implementer, and at what level they want to trust the entity. Using the tree, and different points of trust within that tree, you can get the best of both worlds.
- Q. What about a fully decentralized identity based on protocols not entities?
    - A. Rob: a protocol proves the possession of a secret. Have heard today that we need to limit the amount of secrets in the system. In reality, we need to limit the number of secrets that can leak out. Fundamental to IAM is that “I have a secret which is a private key, and the key proves I am in possession, and you can trust that I am who I am. The problem with that statement is that there needs to be an independent attestation that I am who I am in relation to the private key. Protocols use those private keys to identify people or systems, but without attestation, and your identity federation, anyone can create a self-signed assigned private key and can sign something to a protocol, and can set up a TLS connection, but you have no guarantee that I am who I am, or that I am the man in the middle who is trying to do

a nefarious action. That is the problem saying that you can do this with just protocols.
    - A. Adam. Agree: for that independent party to be able to associate with an actual person gets tricky.
- Q. How do you maintain trust in the federated AM when different nations may have different levels of trust between them?
    - A. Rob: this is the difference between societal trust and confidence. Say an unfriendly nation state is a participant in my trust framework, and they have met all the requirements in the audits, and have mature security protections, and identity and access management, and I know who they are. So at least I know they are my unfriendly nation. My political and societal trusts will say "No. I don't want to accept your information because I know who you are."
- Q. How do you separate the data from the service provider. Currently data belongs to a service provider, how do you break this?
    - A. Rob: service providers have to be part of your trust framework. They need to follow the same rules as the producers and the consumers, and the operators in the system. They cannot work outside the system, and they are an integral part. In ICAO one of the rules we developed was that it is not about the provider, but an assessment of the assets and systems that operate on the data or information, the transactions. So, if they have a set of systems that are in scope for end-to-end operation, then all those systems and assets have to fall under the requirements of the trust framework.
- Q. If using certificates with signing authorities, would you have agreements with the authorities with whose certificates they can sign, and then based on the signature, give appropriate authorization?
    - A. Rob: authorization is an interesting aspect. Typically, certificates are used for identification. The problem with PKI is scalability to authorization. PKI is not agile enough in terms of adding actual attributes to certificates to be able to consider them as attributes for authorization or claims. For me, certificates could be used to identify an entity as having a minimum set of default authorizations, but you would still translate those to a resource manager to a set of default claims based on the identity and origin of identity, and then you would use the authorization system to actually manage their level of authorization on your systems.
- Q. What are your thoughts on authentication/authorization of the transport, which may traverse multiple networks /infrastructure, wired or wireless?
    - A. Rob: If you want to ensure the integrity of information, I believe, and you are in a multi-transport, multi-organizational system, my recommendation is to use message signing, or information signing, which is equivalent in a way to blockchain, except for the lack of the immutable transaction records that you get with the blockchain. If you use signing, I would not care how many different layers or hops between the original consumer exist, and how many handles between service providers. If you use transport layer security, that is only practical if you do it from application to application, say TLS or DTLS. Any time there is a break and inspect or additional providers in between, then you do not have the same assurance. That is why we did

this whole concept with signing messages with the FAA and Eurocontrol, which is also being used by UPP. You can accept a message, but you do not know anything about assurance.

- Q. What do you foresee are the challenges for integrating zero trust architecture into an UM environment?
    - A. Rob: the biggest challenge is that there is no such thing as a way to federate the observable state of a user, or a system. Trustable zero trust is still in its infancy, and really only works in one organization.

Polling Questions, Terrence Lewis

- We do have a polling question related to identity access management. There will be a couple more. Invite the participants to go into those and place your vote.

Blockchain For Aviation Operations in UAM, Richard Walsh, NASA

- Rich Walsh introduced himself; he supports the Convergent Aeronautical Solutions (CAS) at NASA.
- What you see here is "Wicked Works" an idea incubator, that identifier may be a disservice to it.
- We are looking at multiple decades down the road and solving for those. What follows is a concept only. NASA is reviewing our White Paper to determine if this work is to proceed. All comments are sought.
- Blockchain is an emergent technology, primarily business focused. Our novel approach to resolve AAM barriers to scale is theorized to be the backbone for Blockchain applications within the aviation operations domain. We believe it will lead to secure peer to peer communication.
- We seek to establish digital trust of information, and trust in the conveyance of the information, and trust in the use of the conveyed information
- Have used a just released ConOps to prepare.
- Some of the requirements that have a need for secure data are listed. It is our contention that the Blockchain technology can resolve from these requirements.
- The operational areas focus is set by cybersecurity risk in the current setting, otherwise useful data is rendered less trustworthy and not actionable. Broadcast in the open, the data is subject to external manipulation.
- Can Blockchain technology allow for secure peer to peer communication in an aviation operational environment?
- Our approach is to resolve for trust. In this approach, data is converted into a data token which can be transmitted into the open without fear of manipulation. Once the token is unpacked, that data is now useful to the receiving aircraft.
- We propose advancing three experiments that we believe will address two questions:
    - Can tokenization on the Blockchain enable a new level of security, as in ADSB-like data
    - Can that tokenized data then be used in a time-safety critical use case?
- We are looking for what is desirable, what is technically feasible, and commercially viable.

- The barriers include
    - Shift to decentralized third party service providers
    - Increasing autonomous flight with reduced direct human oversight
    - A shift to decentralized third party providers requires a federated operational concept. The Blockchain benefits are aligned
- We want to illustrate the impact that Blockchain technology can have in other aviation operations. One addresses the automation regulations.
- Where are we going? We have submitted the concept document for review. Included were a series of recommendations.
- In summary, the Blockchain operations concept defines a process that creates value. That value may be referred to as proof of location. This novel idea is that parties verify their location relative to each other.
- Quote: "No problem can be solved from the same level of consciousness that created it"

Questions from conferences. Io

- Q. What is the encryption overhead on the ADSB messages?
    - A. Richard Walsh: probably a question for Ken, but we do not know that as yet. Ken Freeman: yes, it would be the size of the token.
- Q. What are eFARs?
    - A. Richard Walsh: they are electronic Federal Aviation Regulations. There is work in the AAM model to incorporate the most important ones so they can be traced to and referred to.
- Q. Is proof-of-location an existing approach or is that idea something you have developed at NASA? Is there any documentation?
    - A. Richard Walsh: we think that it is new, and does not exist. Just coined this. How do we prove where a vehicle is relative to other vehicles?
- Q. What do you think of the computational cost of this may be?
    - A. Richard Walsh/Ken Freeman: this would be defined from the experiments.
- Richard Walsh – I was taking notes from the previous speaker. CAS is about thinking into the future. Imaging trusted authentication of an individual. If an aircraft carries identification with it, certification, tail number, that data is stored as a token that the aircraft takes with it. When the aircraft wants to gain access to the NAS, a handshake of sorts occurs between the aircraft and the NAS, and the entity establishes trust, a Blockchain within that environment.

Immutable Secure Data Exchange and Storage for UAM, Ken Freeman. NASA Ames

- With the UAM Architecture moving towards a decentralized environment with various service providers, we are independent of each other and of the FAA. This work is looking at methodologies for securing the data exchange between those elements as well as the storage associated with it.
- Blockchain is not cryptocurrency. It is the technology behind it, like bitcoin. It provides an immutable distributed ledger that cannot change. It can be distributed across either within or near entities across the airspace structure. It can provide a mechanism of maintaining

security, mostly without a central authority. Basically, you can have distributed aspects that allows, for one, redundancy across the system.

- Within Blockchain you have a list of records, called blocks, and there is a cryptographic hash with a timestamp. By design, once that block is created, it is completely resistant to modification. As new blocks are added, it leverages the hash, and the previous block is part of that process. This helps with the complication of changing a block. It supports that element.
- When we started researching Blockchain a year and a half ago, or two years ago, we needed to make the decision between Ethereum versus Hyperledger Fabric. The first decision was: are we going to a public Blockchain, or permissioned Blockchain.
- Public blockchain is an open protocol. Permissioned blockchains, the transparency is nice, an example would be supply chain tracking. They would not want someone outside of their organization to be part of their blockchain. They will limit the number of nodes from the network.
- We think that UAM should be a permissioned blockchain.
- So why use Blockchain? It is a distributed ledger that we can put our records in. It can be replicated across multiple areas. The different participants will support the maintenance of the blockchain.
- There is a crypto graphic technique that is used to see that the transactions have been added. The ability to validate is relatively simple.
- Immutability is the state of not changing and ensures that once you have something in the block, no one can change that. So, you have an element of trust within that piece of information.
- Consensus is another element of Blockchain. There are different ways to do it, but there needs to be a mechanism to reach consensus. We are currently focusing on a proof-of-work concept, but there are other options.
- When we look at Blockchain from a UAM perspective, we see multiple reasons to use Blockchain. Identity management is one, for the system, pilot, vertiport, or service. Also, It can be used for payment/settlement. It can also be used for Provenance, such as business transactions. It can be used for Data Tracking. See this the most. This applies to UAM for data tracking for flight scheduling, vertiport resources, UAM participant licensing, registration, and UAM system identification.
- This Threat Landscape was mentioned earlier and addresses
    - Man-in-the-middle attack
    - Phishing
    - Malicious attacks
    - Loss of credentials
- Talked about the “Two Generals” Problem
    - One army, two sets of troops, each sitting in a valley, and they cannot see each other, so they have to send messages with the fear of being captured. The message may not reach its destination, or you do not know if the messenger has been captured.

- Applying this to UAM, data could be lost, corrupted, or other. The Two Generals shows unreliable communications and an agreement is never reached. Blockchain can solve this. For UAM, each of the operators would pass their information through the PSU and it would reside on the Blockchain element, then we would put that forward with flight intent, and Blockchain would replicate across the Blockchain, before the next PSU tries to make its intent. This is a good system for parties who do not trust each other.

Questions from conferences. Io

- Q. How would you federate different Blockchains?
  - A. Ken Freeman: If I had multiple Blockchains and a UAM environment, for example, I live in the Bay Area, and Sacramento is to the East of here. Suppose those are operating in different UAM environments and different Blockchains. I see an active area of research in trying to figure out how to interface between those Blockchains. I do not have a good solution for this, but it is an active area of research. It would probably take a piece of software in the middle, between the two.
- Q. Is the Blockchain acting as the trusted central authority?
  - A. Ken Freeman: The use case that we are thinking about at the moment is if vehicles are flying and they have to rely on some function for understanding operational intent, and trusting that function, we do see the Blockchain as a trusted function, but it is decentralized because the element of a Blockchain element may not be the same location for everybody, so it is decentralized. We are seeing it as a trusted entity.
- Q. Is Blockchain verification in series sequence? What if there is a failed block?
  - A. Ken Freeman: think that the verifications are done in series, that's my understanding. There is a centralized bordering service that makes decisions as information is added. There are two sides to it. The intent is added to the Blockchain, and there is something in the middle that makes the decisions, and there is a final part. If there is a failed block, and you try to add to the block chain, and the consensus doesn't come through, then that would simply not be added to the block chain.
- Q. How do you deal with finality and consensus speeds which make the majority of DLTs unsuitable for aviation.
  - A. Ken Freeman: Would say that Blockchain is unsuitable for Detect and Avoid use cases because of the time for it to come to consensus. The work we have done in the past, we are questioning is it even suitable for collecting telemetry directly? We think it is as UML-2, but once we get to UML-4 we are questioning if it is the right solution. But we think that everything else, flight intent, support regulations, things that can support identity management, things that support provenance, those use cases. But anything that would require very quick consensus, getting into the milliseconds, we think that Blockchain is not the right solution. Things that are into the hundred milliseconds, if you can make the decision in that timeframe, we think that Blockchain is viable, and fits the operational intent.

Polling Results

- 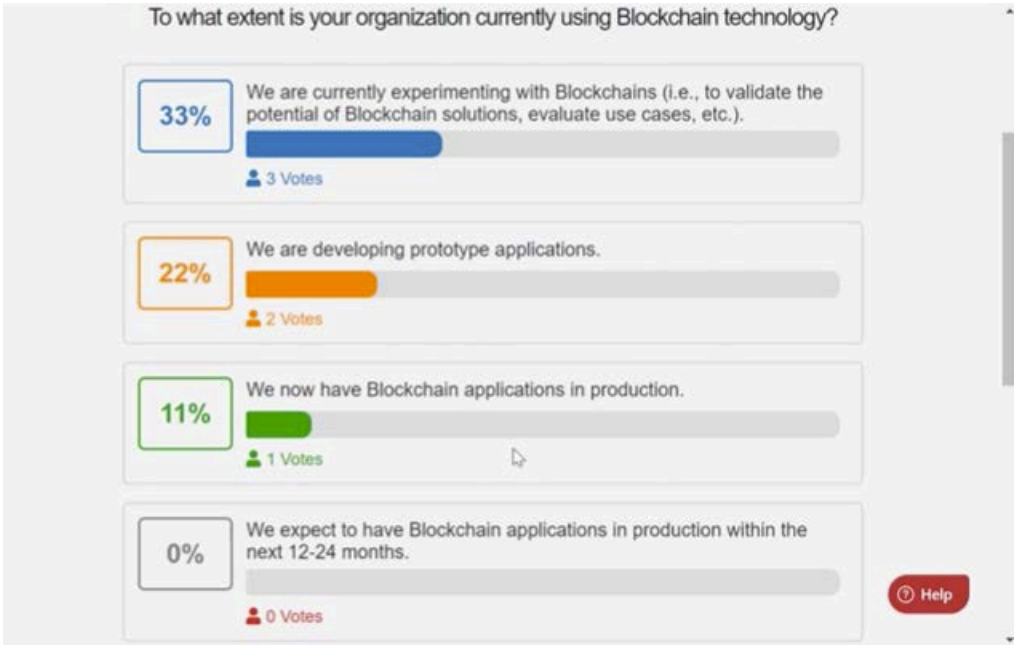

Break – 12:10 to 12:30 pm Pacific

4. DETECT

Model-based Approaches for Cyber Risk Assessment of Space Missions, Arun Viswanathan, JPL

- Arun Viswanathan introduced himself. He leads the Cyber Defense Engineering and Research Group at NASA's Jet Propulsion Laboratory
- The work that we do at JPL is not in the UAM space, but in the space-space. Techniques that we have developed in cyber assessment can be applied to the UAM space.
- Cybersecurity has transitioned from being strictly annoying to having impacts on society. Every year we get new types of attacks against the cyber-physical systems which have real impacts on society. One example is in 2022 on a company named Viasat who owns the KA-SAT satellite. Which brought down internet connection in Ukraine. The impacts are real.
- Our research focuses on Space and Autonomous Systems. We look at attacks on the ground systems, communication links, and spacecraft/autonomous capabilities, and they are all increasing.
- Our group works on securing mission systems, not IT, but cyber defense. We assist missions with cybersecurity across their life cycle phases.
    - The tools we have developed have been deployed across several missions including Mars, Europa, Deep Space Network, and have been in operation for five plus years. Will talk about two of these

    - Another area that we are working on is technology development and fundamental research in cybersecurity. We also work on other non-NASA tasks, such as power grids, oil and gas, DoD.
    - We have been focused on five areas: Intelligent Cyber Defense Technologies for Space, Security of Autonomous and/or Intelligent Systems, Human-machine teaming for cybersecurity (a big interest of mine), Secure System Design and Architecture, and Secure Avionics.
- This is our Challenge slide. Some things here may be interesting to UAM. Cyber for space systems has been hard for us because of the following challenges: culture and mindset; networked, complex and distributed; space research is very collaborative; mission critical systems; legacy systems and components; supply chain risks; long development times; high-value targets.
    - For culture and mindset, we asked why would anybody want to attack us? So, we had to change that mindset before we could start on cybersecurity. This was a huge barrier.
    - Think UAM shares characteristics with "networked, complex, and distributed". You have systems that are owned by different entities, and they have their own security mechanisms and controls, but all have to work together. This makes the problem very challenging.
    - Space research is very collaborative, and we work with other agencies such as the European Space Agency and JAXA. It is hard for us to wall ourselves in. We have to make these collaborations happen.
    - Mission critical systems – this might be shared with UAM spaces. It is not easy to take down things, such as when security patches are required. We have to factor in times when we cannot bring systems down. Helicopters or emergency services—you cannot just shut them down if there is an attack. It takes a different way of thinking how to design these systems.
    - Legacy systems: systems that were launched long ago, software is gone, how do you patch those? People who worked on these systems probably are not at JPL any more. If this is an essential system, it is the weakest link.
    - Supply chain can share this challenge with UAM. Vendors from around the world. Vendors will sub-contract. It is not just hardware. A lot of the code that we use comes from open source, and there are a lot of compromises in the software that open vulnerabilities in critical systems, like at JPL. Most have heard of the JAVA vulnerabilities that occurred in 2022. Cannot imagine the amount of chaos this caused.
    - Then, there are high-value targets. It takes JPL years to develop missions. Designs freeze early in the development, but security is ever-changing. How do you make your mission so it can adapt?
- What we have done to automate cyber risk assessment using model-based approaches
    - It is oftentimes manual, incredibly time consuming, error prone, and too many variables which makes it difficult to do a repeatable cyber assessment and taking into account high-level mission objectives.

    - Our purpose is to capture information about the threat environment, vulnerabilities, identify risks, and evaluate controls. This very difficult. You need good, accurate and concrete information about the system you are trying to assess. Our experience was that it was a months-long process. Information is often not kept in one place, often inside peoples' heads. So, when it comes time to do a risk assessment, it is often not easy to get the information. The process becomes difficult and error prone.
    - We use two main documents, the NIST SP 800-30 and the NIST IR 8179, plus other resources.
- Past challenges show that risk assessments are usually tabletop exercises, using missing, outdated information. People may have left the project, and the time it took was exorbitant.
- We developed an in-house tool, Cyber Analysis and Visualization Environment (CAVE) to do cyber analysis. Referred to the talk earlier by Brian Nolan on MBSE (model-based systems engineering), which we used, but found there were several elements that were an overkill for what we were trying to accomplish, and the learning curve was very steep. We chose a simpler method to model our systems and used a reasoning engine that we built in to do the reasoning.
    - We have created a multi-layered model, all the way from ground models to devices on the networks, to applications, to directories, all the way up to mission functions. You can trace a dependency down to the lowest level, and what is required to support that. These models are assembled automatically. We integrate information such as firewalls, access controls, what software is on the systems.
- We collect data including host data, vulnerability data, network information, mission information from several sources. This process is somewhat manual
- We chose to use a common input format for all data, Excel.csv files. This way a human can review and edit before it gets into the tool
- Once you have all the data, what can you do with it?
    - We can model it all the way up to the spacecraft, but right now just modeling at the ground. The following questions would be very hard to answer if you have a complicated distributed environment.
    - The reasoning questions can include: which ports can two servers communicate? What mounted directories can a server read? Are there any critical vulnerabilities on servers that can run a mission-critical application? Which systems have a vulnerability with a downloadable exploit? Can an adversary access a critical mission resource from the internet?
    - You can, say ask the model to show me all the systems that are vulnerable, and it is impossible to answer If all you have in the current environment are a bunch of Excel sheets, or diagrams. The ability to model and reason has been game changing.
    - Also, on one of the missions, we asked whether an adversary can access a critical mission resource, using wifi? Imagined the answer would be "no", but that was not the case. We found an open port that someone had forgotten to close. Detecting this would not have been possible if it had not been automated.

- We asked CAVE to find all paths on known vulnerable ports to servers that have access to the command directory; find all servers with vulnerabilities with a downloadable exploit that compromise data integrity; find all servers with critical vulnerabilities that have exploits available and run mission-critical software. If an attacker gains access to the command directory, which is the crown jewels for any mission, they can potentially force their own commands.
    - Threat analyses were then conducted by the mission cybersecurity team. Science data and telemetry are also one of the crown jewels for us.
    - If attackers can access servers with critical vulnerabilities, they can access critical mission systems, which could potentially compromise the whole mission.
    - One analysis showed that one of those systems, which was the most vulnerable server had a lot of known exploits. This is the kind of insight that can be scanned and is immediately actionable, which was previously not possible. Showing this to any level of project management structure speaks immediately to them.
- Summary
    - Models consolidated and siloed "tribal knowledge". The edges between the nodes had different semantics. Connecting a host to a host, the edge has a different meaning than connecting a file to a host. The reasoning takes care of those semantics.
    - Model-based reasoning has dramatically reduced the time to identify security weaknesses from months to minutes
    - Model-based risk assessments are scalable, repeatable, and accurate
    - Feel free to contact us if you want more information or a demo on the tool.
- Arun Viswanathan shared a demo from his computer browser using a use case that applied CAVE to a mission project plan and perform a risk assessment. This is a representation of a mission system, cleared for sharing.
    - Analyses are customizable. This was demonstrated

Questions from conferences. Io

Q. Are these risk assessments mostly performed on systems that are further along in the design process? Or have you been able to integrate earlier into a project's lifecycle to help inform secure design decisions?

- A. Arun Viswanathan. We began using risk assessment systems that were already operational. Now we are at a stage where we are informing the design at earlier phases of the design, and we can now do "what if" scenarios. So can borrow an existing design and start customizing it, and then you can do "what if" using this tool

In-time National Airspace Security Event, Paul Hoyt Nelson, NASA Glenn

- Thank you all for coming. I am Paul Hoyt Nelson, the Senior Cybersecurity Advisor for the Aeronautics Research Mission
    - Am responsible for ensuring that ARMD's work both secure and securable when it gets moved out of the Agency, a broad set of responsibilities.
    - For the AAM/UAM work, I own the security requirements work that is going to be matured in 2032. It fits in with all of the things we have been talking about.

    - Want to show you the work in the Safety Management System and will show you the corollaries to that and how we are starting to tie them together in research. The key point is that this is research, not production or operational systems at this point
- A quick note on the timeline, moving along the Epochs, starting in the late 40s when we really started looking at Aviation Safety, we are exiting Epoch 3, where we defined trajectory-based air traffic management and systems, and entering into Epoch 4 collaborative, connective phase where data and services manage the way we drive the airspace.
- Will go through the In-Time Safety Management System (IASMS). A chart was shown that shows an overview of the whole collaborative system environment. It includes all of our organizations, the FAA, aviation, airlines, the types of systems (UAM, high altitude, commercial transport) all of the things that are in the commercial aviation system today.
- The safety part shows that Safety Policy and Safety Promotion are the way we have done safety for ever in the National Airspace System. In Policy, we have broad safety objectives, and responsibility and accountability. In the Promotion side, it's training, information flows, and developing a safety culture. In the middle of these is the In-Time System-Wide Safety Assurance (ISSA) capability that is implemented in the IASMS. It is built on services, functions, and capabilities. We have processes that we monitor>assess>mitigate—iteratively all the time.
- Now we look at the service/function/capabilities from more of an operational point of view. We have "manage" (down draft/thunderstorms/airport shutdown). We have "identify" (the unknown risk parts). We have "inform" (the analysis going forward). In-Time is a critical point. It was "Real-time" at one point. We wanted to identify things "In-Time" to do something about things. In some cases that could be milliseconds, such as a separation issue. In other cases it could be months or years if it is a training or design issue.
- If we look at the Monitor/Assess/Mitigate, we have different data sets coming in, flight plans, meteorology, configurations, and they flow into a data fusion model where we monitor the state, detect precursors, and lead indicators. That flows into a prognostic and predictable analytical engine that uses advanced algorithms to assess the operational data and understand what is happening. That ends up being an alerting phase and/or a decision support phase. Again, an alert might not be the right thing if it is an airframe issue that is going to take years to address.
    - The mitigation response is time-dependent mitigation action, procedure-based (augmented using the decision-support tools). That flows right back into the overall system. That is your circular Monitor/Assess/Mitigate system.
- It all starts with data, which is why we are here with the cybersecurity part. This data comes from a lot of different places, such as ANSP, Operators, Vehicles, Supplemental Data Service Provider, System Wide Information Management and Flight Information Management System, and other sources. There is a wide variety of data that fits into this. All of these are targets for safety, but also for security.
- The IASMS has developed UAS, UTM, a service-oriented architecture. We have the vehicle SFCs, the providers of UAM services and operators, and the SDSPs and almost anything you

can imagine. So together this is a high-level view of the Service Oriented Architecture. Everybody is a potential service provider and a potential service consumer.
    - A quick trip down into IASMS and how it helps with cybersecurity. This work started in 2018.
        - This is operational and technology data, not business, and is directed at physical vehicles, things that are flying or things on the ground, in the real world.
        - Safety hazards represent a good target for bad actors that have some digital cyber capability.
        - The interrelationships are well understood and predictable
        - The Indicators of Compromise (IoC) data would not be identifiable using normal IT security methodologies
        - The OT data describes the real world. We will use this to identify bad actors
- A research question is: can in-time warning of cybersecurity incidents be predicted from ISSA data feeds? This is what a System-Wide Safety team is looking into this
- A couple of examples: ADS-B identifies where aircraft are in the National Airspace System. It is a transceiver on an aircraft. It uses GPS to understand where it is, and it transmits the information to the ground and to other aircraft. It is highly dependent upon GPS to be accurate. There are known attacks that cause GPS to be highly inaccurate. Not jamming, but spoofing, where you are changing the understanding of where the GPS actually is.
    - We know we can deviate GPS by many miles, 50 or 60 miles. That could put aircraft in the position of collisions, so we need to understand what is happening. Analyzing the ADS-B data over time means we can develop a pattern of vectors and speeds. We know planes fly at certain speeds, based on the type of aircraft. So, a Cessna 172 is not going to be at 300 knots, but probably tops out at 125. We know that a Boeing 737 is not going to be flying at 125 at altitude. It would be moving much faster. All of this pre-knowledge has to adhere to the physics model. We can build this in near real time, and it allows us to evaluate whether or not the plane as it is supposed to behave is behaving the way it should be. This is a way of identifying that the GPS data that the aircraft believes is actually correct. So, we know maximum/minimum/climb/descent rates, and we can identify if there is an indicator of compromise to the vehicle to its GPS. Let that sink in. Physics tells us if we are being attacked.
- Weather data is kind of the same category of the physics model.
    - We all know this weather data exists: forecasts/analysis/observations. Weather changes quickly and slowly and depends on the type of weather in the environment it is forecast to behave in. If no hurricane is predicted in the forecast, then you are not going to see one.
    - Vehicles and air crews use this data to navigate, and they also provide observations. That builds a more complete and precise picture of what weather is supposed to look like.
- An attack on the data service, which we described some years back with some other NASA work, because the weather service was not providing any kind of data integrity with the data feeds. So, the attack on the data service could result in loss of efficiency, an attack on the National Airspace System, or an attack on the airline. But it is a denial of service.

    - Identifying a data attack is possible by fusing that data that we know together: air crew observations, vehicle performance data, ground observations, vehicle sensor data. All of those build a picture of what the weather is, versus what we are being told what the weather is.
    - That allows us to go back to the assess/monitor/mitigate approach, and to build an identification of a data attack.
- Conceptually, going back to that view of all of the data types in the aviation system that IASMS is using, as we look at the data types and fusing them, we have the description of the physical part of the system that we can validate, and then we know we have potentially got something we can look at.
- The IASMS is going to be doing its own analysis, built along the model shared in this chart. It takes the information exchange from the NAS/airspace management/fleet management/vehicle management and the operational systems. They all flow together into this entire environment that we are now familiar with: Air Traffic Control, USS, PSUs, and that goes into the IASMS. The supplemental data service providers and others all flow into the IASMS. There are currently different views on what that piece will look like. It allows us to look at all of those pieces coming together and potentially identify a compromise of the entire system. Take, for example, firefighting. Maybe a leading indicator for a larger NAS attack, because this is attacks against the real world, they don't necessarily cascade the way they do in the virtual world. So, we have potential there, looking at the way the IASMS is doing its own analysis, and comparing that with the data that is being provided.
- A description of research work that is going on.
    - The research cycle is essentially on an annual basis, based on a systems analysis of the IASMS and data feeds going into it to:
        - Develop Use cases
        - Develop gap analysis on use case vs. reality
        - Develop cyber security concept proposals
        - Modeling and simulating in a system model
    - Looking at weather again, we look at weather, and the analysis of that coming in to develop a couple of use cases. Due to gap analysis, we cannot rely on the integrity of our weather data. Our concept proposal is to go back to NOAA and have them start to provide integrity data for that.
    - We believe that over the years we will have far more complex use cases and cybersecurity concepts.
    - Our goal is to develop a security analog to the In-Time Aviation Safety Management System that is not separate, but piggybacked onto that technology.

Questions from conferences. io

- Q. Do you think there's enough discrete weather data right now to ensure safety, particularly for micro-weather (urban)?
    - A. Paul Hoyt Nelson: No. The first problem is we don't really have the monitoring systema in place in any kind of consistent basis. We could take performance data off the vehicles. We have had UAM vehicles fly backwards because they did not have enough propulsion against the wind. The other problem is the integrity of the

provider of the data. We do not have any system in place for validated data. We do not have anything in the system to consume that data and provide any meaningful feedback. It is going to be very necessary for operations in urban areas.

- Q. What type of data do you expect to collect from an air vehicle OEM?
    - A. Paul Hoyt Nelson: We bounced around a lot of this. One of the concepts had to do with engine Full Authority Digital Engine Control (FADEC) data and whether it would be useful. Does it make sense to flow that information out? Maybe. If there is a failure mode in the engine, the air crew is going to find out about it and take action. If the system is telling them sooner, then the aircraft can respond based on that possibility. However, it is different with battery-powered vehicles, that do not have high levels of endurance. With a jet, I can loiter, with a battery-powered, I cannot loiter, I will need to set down in the next 15–20 minutes. That flowing into the system could be very useful.
- Q. One controversial area in an aircraft is organophosphate poisoning. Could multiple data sets offer safety insights?
    - A. Paul Hoyt Nelson: I do not know. There are ongoing discussions on contrails in general, and what can be done to minimize the impact on the climate. They are complex, and not well understood as to what they contain, and how long they last. Referenced 9/11 when contrails vastly disappeared, and weather changed.

Poll Results:

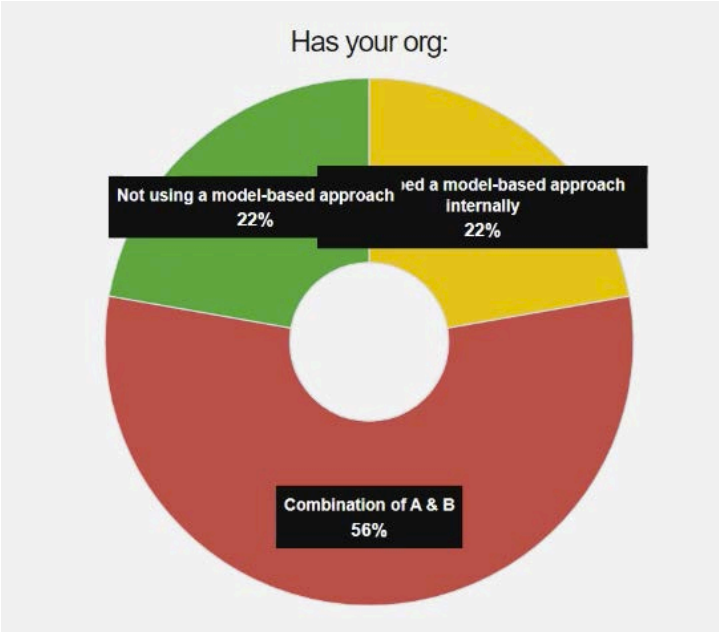

The Workshop was adjourned at 2:05 pm Pacific with closing remarks by Terrence Lewis

Summary

This workshop outbrief report of the NASA's Secured Airspace for UAM Workshop discussion outlines the potential recommendations of capabilities to improve the understanding of UAM cybersecurity. It should be noted that these topic discussions and outcomes are not exhaustive, nor do they represent a consensus among participants. The Secure Airspace Technology Group at NASA continues to work with and learn from various members of the UAM and cybersecurity community and look forward to future chances for collaboration.

Appendix A | NASA's Secured Airspace for UAM Virtual Workshop
Tuesday, November 1, 2022

Topics – Schedule 8am-2pm PT (11am-5pm ET)

Topics	Time	Notes
• Welcome	8:00 a.m. – 8:05 a.m.	Kevin Witzberger – NASA ARC
• Goals of the Workshop	8:05 a.m. - 8:15 a.m.	Terrence Lewis – NASA ARC
UAM		
• What is UAM ?	8:15 a.m. – 8:45 a.m.	Annie Cheng – NASA ARC
Identify		
• Challenges in Securing UAM Operations	8:45 a.m. – 9:15 a.m.	Kenneth Freeman – NASA ARC
Protect		
• Principles of a Trustworthy System	9:15 a.m. – 9:45 a.m.	Gano Chatterji – Crown
• Security Architecture	9:45 a.m. – 10:15 a.m.	Bryan Nolan
• Break	10:15 a.m. – 10:45 a.m.	Break
• Identity Access Management Panel	10:45 a.m. – 11:30 a.m.	Adam Monsalve – CNA & Robert Segers - FAA
• Poll #1		
• Blockchain for Aviation in UAM	11:30 a.m. – 12:15 p.m.	Richard Walsh - NASA ARC
• Poll #2		
• Immutable Secure Data Exchange and Storage for UAM	12:15 p.m. – 12:45 p.m.	Kenneth Freeman – NASA ARC
• Break	12:45 p.m. – 1:00 p.m.	Break
Detect		
• In-Time National Airspace Security Event	1:00 p.m. – 1:30 p.m.	Paul Nelson – NASA GRC
• Model-based Approaches for Cyber Risk Assessment of Space Missions	1:30 p.m. – 2:00 p.m.	Arun Vishwanathan – NASA JPL
• Poll #3		

1